INTELLIGENT CAI

BOLT BERANEK AND NEWMAN, INCORPORATED

PREPARED FOR
OFFICE OF NAVAL RESEARCH
ADVANCED RESEARCH PROJECTS AGENCY

OCTOBER 1975

KEEP UP TO DATE

Between the time you ordered this report—which is only one of the hundreds of thousands in the NTIS information collection available to you—and the time you are reading this message, several *new* reports relevant to your interests probably have entered the collection.

Subscribe to the **Weekly Government Abstracts** series that will bring you summaries of new reports *as soon as they are received by NTIS* from the originators of the research. The WGA's are an NTIS weekly newsletter service covering the most recent research findings in 25 areas of industrial, technological, and sociological interest—invaluable information for executives and professionals who must keep up to date.

The executive and professional information service provided by NTIS in the **Weekly Government Abstracts** newsletters will give you thorough and comprehensive coverage of government-conducted or sponsored research activities. And you'll get this important information within two weeks of the time it's released by originating agencies.

WGA newsletters are computer produced and electronically photocomposed to slash the time gap between the release of a report and its availability. You can learn about technical innovations immediately—and use them in the most meaningful and productive ways possible for your organization. Please request NTIS-PR-205/PCW for more information.

The weekly newsletter series will keep you current. But *learn what you have missed in the past* by ordering a computer **NTISearch** of all the research reports in your area of interest, dating as far back as 1964, if you wish. Please request NTIS-PR-186/FCN for more information.

WRITE: Managing Editor
5285 Port Royal Road
Springfield, VA 22161

Keep Up To Date With SRIM

SRIM (Selected Research in Microfiche) provides you with regular, automatic distribution of the complete texts of NTIS research reports *only* in the subject areas you select. SRIM covers almost all Government research reports by subject area and/or the originating Federal or local government agency. You may subscribe by any category or subcategory of our WGA (**Weekly Government Abstracts**) or **Government Reports Announcements and Index** categories, or to the reports issued by a particular agency such as the Department of Defense, Federal Energy Administration, or Environmental Protection Agency. Other options that will give you greater selectivity are available on request.

The cost of SRIM service is only 45¢ domestic (60¢ foreign) for each complete microfiched report. Your SRIM service begins as soon as your order is received and processed and you will receive biweekly shipments thereafter. If you wish, your service will be backdated to furnish you microfiche of reports issued earlier.

Because of contractual arrangements with several Special Technology Groups, not all NTIS reports are distributed in the SRIM program. You will receive a notice in your microfiche shipments identifying the exceptionally priced reports not available through SRIM.

A deposit account with NTIS is required before this service can be initiated. If you have specific questions concerning this service, please call (703) 451-1558, or write NTIS, attention SRIM Product Manager.

This information product distributed by

NTIS U.S. DEPARTMENT OF COMMERCE
National Technical Information Service
5285 Port Royal Road
Springfield, Virginia 22161

311082

BOLT BERANEK AND NEWMAN INC

CONSULTING · DEVELOPMENT · RESEARCH

BBN Report No. 3181

INTELLIGENT CAI

Final Report for Contract No. N00014-71-C-0228

Allan Collins
Mario C. Grignetti

October 1975

PRICES SUBJECT TO CHANGE

CAMBRIDGE   WASHINGTON, D.C.   CHICAGO   HOUSTON   LOS ANGELES   SAN FRANCISCO

REPORT DOCUMENTATION PAGE		READ INSTRUCTIONS BEFORE COMPLETING FORM
1. REPORT NUMBER Technical Report No. 6	2. GOVT ACCESSION NO.	3. RECIPIENT'S CATALOG NUMBER
4. TITLE (and Subtitle) Intelligent CAI		5. TYPE OF REPORT & PERIOD COVERED Final Report (1 March 71-31 Aug. 75)
		6. PERFORMING ORG. REPORT NUMBER BBN Report No. 3181
7. AUTHOR(S) Allan Collins, Mario C. Grignetti		8. CONTRACT OR GRANT NUMBER(S) No. N00014-71-0228
9. PERFORMING ORGANIZATION NAME AND ADDRESS Bolt Beranek and Newman Inc. 50 Moulton Street, Cambridge, Mass.		10. PROGRAM ELEMENT PROJECT TASK AREA & WORK UNIT NUMBERS 61153N RR042-06;PR042-06-01 NR154-330
11. CONTROLLING OFFICE NAME AND ADDRESS Personnel and Training Research Programs Office of Naval Research (Code 458) Arlington, VA 22217		12. REPORT DATE October 1975
		13. NUMBER OF PAGES 47
14. MONITORING AGENCY NAME & ADDRESS(if different from Controlling Office)		15. SECURITY CLASS. (of this report) Unclassified
		15a. DECLASSIFICATION/DOWNGRADING SCHEDULE

16. DISTRIBUTION STATEMENT (of this Report)

Approved for public release; distribution unlimited.

17. DISTRIBUTION STATEMENT (of the abstract entered in Block 20, if different from Report)

18. SUPPLEMENTARY NOTES

19. KEY WORDS (Continue on reverse side if necessary and identify by block number)

Education, Tutoring, Computer-assisted instruction, Learning, Generative CAI, Teaching

20. ABSTRACT (Continue on reverse side if necessary and identify by block number)

This paper describes the capabilities now available for building intelligent, tutorial CAI systems as exemplified by several systems including Tutor-SCHOLAR, Map-SCHOLAR, NLS-SCHOLAR and SOPHIE. The systems illustrate how a variety of sophisticated techniques can be used for tutoring different kinds of knowledge by carrying on dialogues in natural language. The systems have been developed to explore how to provide each student with his own personal, expert tutor.

Intelligent CAI

Allan Collins
Mario C. Grignetti

Bolt Beranek and Newman Inc.

Cambridge, Massachusetts 02138

Annotated Bibliography of Papers Prepared for this Project

Carbonell, J. R. and Collins, A. M. Natural semantics in
    artificial intelligence. In Proceedings of the Third
    International Joint Conference on Artificial Intelligence,
    Stanford University, 1973, 344-351. Reprinted in the
    American Journal of Computational Linguistics, 1, Mfc 3, 1974.
    (Partial support)

    This paper discusses human semantic knowledge and
    processing in terms of the SCHOLAR system. In one
    major section we discuss the imprecision, the incom-
    pleteness, the open-endedness, and the uncertainty of
    people's knowledge. In the other major section
    we discuss strategies people use to make different
    types of deductive, negative, and functional inferences,
    and the way uncertainties combine in these inferences.

Collins, A. M., Passafiume, J. J., Gould, L., and Carbonell, J. G.
    Improving interactive capabilities in computer-assisted
    instruction. BBN Report No. 2631, 1973.

    This report describes the development of interactive
    capabilities in the SCHOLAR CAI system centering in three
    main areas: (1) implementation of two presentation
    strategies in SCHOLAR (Tutorial mode and Block-Test mode)
    and a comparative evaluation of these two modes using
    high-school students as subjects; (2) initial study based
    on analysis of tutorial dialogues of how to teach

1

procedural knowledge interactively within SCHOLAR, and
(3) addition of a module for teaching geography using the
map display and related question-answering facilities
recently added to SCHOLAR.

Collins, A. M., Warnock, E. H., and Passafiume, J. J. Analysis
   and synthesis of tutorial dialogues. In G. Bower (Ed.),
   The psychology of learning and motivation, Vol. 9.
   New York: Academic Press, 1975.

In this paper we attempt to analyze the strategies by
which tutors adapt their teaching to individual students,
so that we can synthesize these strategies in the
SCHOLAR CAI system. To find out what strategies
tutors use, we tape-recorded dialogues between various
tutors and students on the topic of South American
geography. Because SCHOLAR is a well-defined program,
it is possible to analyze such ill-defined naturalistic
data in precise terms, with respect to the structure
and processing of information in SCHOLAR. We analyzed
the dialogues concentrating on one aspect at a time.
Based on our analyses, we propose in this paper several
hypotheses about how the tutor relates his teaching to
the individual student. We show how in modified form
we have implemented some of these strategies in SCHOLAR.
We further argue that the analytical method employed here
could be extended to a wide range of conversational
situations. This method (Dialogue Analysis) would permit
psychologists to study questions about the interactive
aspects of human processing that cannot even be considered
with traditional laboratory methods.

Collins, A. M. Comparison of two teaching strategies in computer-assisted instruction. BBN Report No. 2885, 1974. Submitted to Instructional Science.

Three experiments were run using the SCHOLAR CAI system to teach geography to high-school students. The experiments compared a method of teaching derived from analysis of human tutors (Tutorial mode) vs. a method derived from programmed instruction (Block-Test mode). In the three experiments, Block-Test mode was systematically converged toward Tutorial mode in order to pinpoint what aspects of teaching strategy affected students' learning. Tutorial mode was significantly more effective in the first two experiments, and nonsignificantly in the third. The results indicated that the major factor affecting students' learning was the strategy that tutors use of reviewing the material in greater depth on a second pass. Allowing the students to ask questions, and the tutorial strategy for relating new material to the students' previous knowledge contributed only a small amount to the differences found in the first two experiments.

Grignetti, M. C., Hausmann, C., and Gould, L. An "intelligent" on-line assistant and tutor--NLS-SCHOLAR. In Proceedings of the National Computer Conference, San Diego, California, 1975, 775-781. (Partial support)

NLS-SCHOLAR is a prototype system that uses Artificial Intelligence techniques to teach computer-naive people how to use a powerful and complex editor. It represents

a new kind of Computer Assisted Instruction (CAI) system that integrates systematic teaching with actual practice, i.e., one which can keep the user under tutorial supervision while allowing him to try out what he learns on the system he is learning about.

NLS-SCHOLAR can also be used as an on-line help system outside the tutorial environment, in the course of a user's actual work. This capability of combining on-line assistance with training is an extension of the traditional notion of CAI.

The techniques used in NLS-SCHOLAR are general and can be applied to a wide variety of computer-related activities.

Collins, A. M., Warnock, E. H., Aiello, N., and Miller, M. W. Reasoning from incomplete knowledge. In D. G. Bobrow and A. M. Collins (Eds.), Representation and understanding. New York: Academic Press, 1975.

The paper describes how people use a variety of plausible, but uncertain, inferences to answer questions about which their knowledge is incomplete. This kind of reasoning is described in terms of how it is being implemented in the SCHOLAR CAI system. The paper also shows how people can be taught to reason in this way, using a Socratic tutorial method implemented in a system like SCHOLAR.

Collins, A.   Education and understanding.   In D. Klahr (Ed.)
Cognition and instruction. Hillsdale, N. J.: Erlbaum
Associates, 1975.

> This chapter comments on chapters by Just and Carpenter
> and by Simon and Hayes on teaching understanding skills.
> The chapter argues that the most important aspect of
> understanding is how people use their knowledge about
> the world to fill in the information the text assumes.
> Thus there can be no easy way to educate people to
> understand, because they need to be taught both a huge
> amount of world knowledge and the understanding skills
> to use that knowledge effectively in reading.

Collins, A. and Grignetti, M. C.   Intelligent CAI.
BBN Report No. 3181, 1975. To be submitted to Science.

> This paper describes the capabilities now available for
> building intelligent, tutorial CAI systems as exemplified
> by several systems including Tutor-SCHOLAR, Map-SCHOLAR,
> NLS-SCHOLAR and SOPHIE. The systems illustrate how a
> variety of sophisticated techniques can be used for
> tutoring different kinds of knowledge by carrying on
> dialogues in natural language. The systems have been
> developed to explore how to provide each student with
> his own personal, expert tutor.

Collins, A., Pew, R. W., and Adams, M.   The effectiveness
of an interactive map display in tutoring geography. In
preparation.

> This paper will describe the Map-SCHOLAR system and an
> experiment that compares how well students learn in
> Tutorial Mode, using (a) the interactive map display
> of Map-SCHOLAR (b) a static labeled map, and (c) an
> unlabeled map. The paper will also show how a new
> method called _backtrace analysis_ can be used to
> pinpoint the effectiveness of different aspects of
> the tutoring strategy and the map system used in the
> experiment.

Intelligent CAI

Allan Collins

Mario C. Grignetti

Bolt Beranek & Newman Inc.

Cambridge, Mass. 02138

INTRODUCTION

If computerized instruction is ever to have a large impact on education, computer-assisted instruction (CAI) systems must have the flexibility and skill of a human teacher. In developing the SCHOLAR CAI system Carbonell (1) took a first step toward an intelligent tutorial CAI system. In SCHOLAR, knowledge was not stored as text, but in an interrelated network of facts and concepts, so that the knowledge could be used in a variety of ways. In short the attempt was to structure information like a human knowledge, so that the program could use its knowledge as flexibly as a human tutor does.

In this paper we will discuss the structure of the SCHOLAR system, some of the ways that the potential for intelligent CAI has been realized in current systems, and finally what is possible in the near future toward building intelligent tutorial systems.

The Context of SCHOLAR in CAI

Prior to Carbonell's SCHOLAR program, CAI had proceeded along several lines. Bryan (2) distinguished three broad categories. In the first, ad-hoc CAI, the student is given full control of the computer with a simple programming language and perhaps a series of tasks to perform. LOGO (3) provides one of the most interesting educational environments of this kind and indeed children learn some important cognitive skills in working with LOGO. The second category is games and simulation, where the student learns indirectly while participating in the game or simulation. The Plato

system's "How the west was won?" (4) is an excellent example of such a system where children learn the arithmetic operations in playing a variant of "Chutes and Ladders." Both these forms of CAI are highly interactive, but they are limited as teaching methods to certain kinds of knowledge.

The third category Bryan called controlled learning. Most programs in this category specify the possible sequences through a program, where different branches are taken depending on the student's responses to questions or problems. The sequence a student follows is usually deterministic, with a branch for each anticipated class of responses by the student (sometimes based on a keyword he might give). Some ingenious programs can be written in this way, such as the Socratic system (5) or the chemistry programs in the Plato System (6), but there are some inherent limitations to this approach. The student can not use natural language in his responses, and cannot ask any but specifically anticipated questions. The teacher has a considerable burden in the preparation of questions, answers, keywords, and branchings. From a system's point of view, the system has no real initiative or decision power of its own, nor any knowledge that is available other than at fixed points in the sequence.

The rigidity of this latter approach led to the development of "generative CAI" (7). The first CAI programs were based on mathematics and other well-defined subjects, where problems could be generated and answers evaluated by "algorithmic manipulation". This

was a major step toward representing a program's knowledge so that it could be used in a variety of ways. But the knowledge was limited to algorithmic domains.

Carbonell's contribution was to apply the idea of generative CAI to ill-defined verbal knowledge, in particular the domain of geography. To do this he used a semantic network (8) for representing knowledge about the world. In a semantic network, information is stored about concepts under different entries, much as in an encyclopedia, but according to a well-defined format. Every concept used to describe a given concept can itself be described elsewhere in the network. Thus the program can know about every concept it uses.

When information is represented in the general format a semantic network, then the structure of the network can be used by different procedures to formulate questions to the student, evaluate his answers, answer his questions, make inferences and computations, select new topics for discussion, etc. Freeing the information from the specific use it is to be put makes it available for use in a variety of ways. In this paper we will try to show some of the power for tutorial interaction that can be obtained by such an approach to CAI.

The Original SCHOLAR Program

Figure 1 presents a short protocol with SCHOLAR, showing some of the original system's capabilities. SCHOLAR selected topics on a

WHAT ARE THE PRINCIPAL CITIES OF BRAZIL?

RIO AND MAYBE BUENOS AIRES OR SAN PAULO

RIO DE JANEIRO AND SAO PAULO ARE CORRECT. YOU MISSPELLED
SAO PAULO. BUENOS AIRES IS IN ARGENTINA, NOT IN BRAZIL.
THE OTHER PRINCIPAL CITIES OF BRAZIL ARE BRASILIA, PORTO
ALEGRE, SALVADOR, RECIFE, SANTOS, MANAOS, AND BELEM.

TRUE OR FALSE? THE LANGUAGE IN BRAZIL IS FRENCH.

Q/A

DO YOU WANT TO CHANGE THE GENERAL MODE OF THE INTERACTION TO
QUESTION-ANSWERING MODE? PLEASE ANSWER Y OR N.

Y

OK, MR. JONES. LET'S GO AHEAD.

WHERE IS BRASILIA?

BRASILIA IS IN CENTRAL BRAZIL. IT IS IN THE BRAZILIAN
HIGHLANDS.

TELL ME MORE ABOUT THE LOCATION OF BRASILIA?

THE LATITUDE OF BRASILIA IS -16. THE LONGITUDE IS -48.

WHAT IS LATITUDE?

A LATITUDE IS AN ANGULAR DISTANCE FROM THE EQUATOR. IT CAN
BE CONSIDERED A PART OF A LOCATION.

*WHAT IS THE MOST IMPORTANT DIFFERENCE BETWEEN BRASILIA AND
RIO*

THE PRIMARY DIFFERENCE IS: BRASILIA IS IN CENTRAL BRAZIL.
BUT RIO DE JANEIRO IS IN SOUTHEASTERN BRAZIL.

IS THE CLIMATE OF SAO PAULO TROPICAL

NO. THE CLIMATE OF SAO PAULO IS TEMPERATE.

Fig. 1 Dialogue with the original SCHOLAR

random basis, weighted for importance. When it printed an asterisk, the student could respond to the question by SCHOLAR, direct a question to SCHOLAR, or issue a command such as halting or changing the mode of interaction. The student returned control to SCHOLAR by typing an asterisk and carriage return.

SCHOLAR's questions were constructed from the information in the semantic network, and were not prestored. The student's answers were evaluated by different procedures with respect to the information in the network, deciding whether the particular answer was correct, wrong, partly correct, or approximately correct. Student's questions were analyzed syntactically to determine what information was sought, and the semantic network searched to find the information. Different computations and inferences were applied, if the information was not stored directly. All these operations were carried out by procedures which operated independently of the specific information that was involved.

In Fig. 1 the questions by SCHOLAR illustrate different kinds of questions that could be generated. After the second question by SCHOLAR, the student changed the mode of interaction from mixed-initiative mode, where SCHOLAR asked him questions, to question-answering mode, where SCHOLAR waited for questions from the student. Other modes described below have since been added to SCHOLAR. Because the student can control the way he interacts with the system, he can choose the mode of interaction that he finds most effective. This is one of the important ways such a system

personalizes instruction.

The student then asked a series of questions to clarify and extend the information given to him about the cities of Brazil. When the student wants more information about something such as Brazilia, he can ask specifically what he wants to know. When he doesn't understand a word, such as latitude, he can have it explained. In this way the knowledge taught can be geared to the individual student's background, so as not to repeat what he already knows or go over his head. This is important to maintaining a student's motivation to learn.

The two questions about the location of Brasilia illustrate how a tutorial system can avoid overloading the student with too much information at one time. Each piece of information in the network is tagged to indicate its relative importance. The program gives only the most important information at any time, but the student can always ask for more information if he wants it.

The questions about the most important difference between Brasilia and Rio and about the climate of Sao Paulo illustrate the ability to use a semantic network to make appropriate computations and inferences. In the first case there is a procedure for comparing two things to find their similarities and/or differences. Each property of the two things is compared in the order of importance. Here the most important property on which the two things differ is found, and given as an answer to the student. The second case illustrates a combination of two inferences, a deduction

and a contradiction. Nothing about climate is stored with Sao Paulo, but Sao Paulo is in the Brazilian Highlands which has a temperate climate. By comparing tropical and temperate, SCHOLAR finds there is a contradiction and concludes the answer is "no". There are a large number of such inferential strategies that humans use, and only some of the more common ones have been implemented in SCHOLAR. But information in SCHOLAR is structured in such a way that it is possible to specify content-independent procedures to carry out different inferences.

This summarizes the major contributions of the original SCHOLAR system. There were also several severe limitations to the original SCHOLAR. First, the information in the program was restricted to static, verbal facts about geography, which are not very interesting in themselves. Second, the program was quite restricted in its ability to understand student answers and questions because of its limited language processing capability. Third, and perhaps most important there was no teaching strategy; the program merely generated questions randomly or answered student questions. We will try to show how later systems have overcome some of the limitations of the original SCHOLAR and at the same time exploited further its potential for tutorial interaction.

TEACHING STRATEGY

Analysis of Human Tutoring

In the original SCHOLAR there was no teaching strategy, but SCHOLAR's structure made it possible to model the way human tutors interact with students. By collecting tape recordings of different tutors teaching the same kind of information as SCHOLAR, it was possible to analyze how tutors adapt their teaching to the individual student (9). There were four crucial aspects of their tutoring strategy, that were subsequently modelled in SCHOLAR. They were: (a) the way tutors select topics, (b) the way they interweave questions and presentation, (c) their reviewing, and (d) their error correction strategy.

The topic selection strategy used by tutors produces a structure of topics and subtopics like an outline for a course. For example, the tutor might start off with a question like "Do you know any geographical features of South America?" If the student gives Cape Horn, for example, then the tutor would discuss Cape Horn for a while, including perhaps the Straits of Magellan as a subtopic. After covering the most important information about Cape Horn, the tutor would then ask about other geographical features, like the Amazon or the Andes. Each of these would be discussed briefly until the major geographical features are covered, at which point the tutor would pick a new topic such as regions or countries. Thus, the topics and subtopics form a nested outline structure, with the

tutor probing a little way into each subtopic, and then popping up to the previous topic when the important information is exhausted. The better the tutor, the more structure there is to the discussion.

The way the tutor interweaves questioning and presentation is the essence of how the tutor relates his teaching to the individual student. The dialogues showed that the tutors' questions occur at the top-level and beginning topics in the outline. This is because the tutor starts out asking questions to find out what the student already knows, and then presents new material that is related to the student's previous knowledge. The object seems to be to tie as much information as the student can assimilate into the structure of his previous knowledge (10).

Another important aspect of the tutorial strategy is reviewing. In the dialogues the better tutors went over the material on a second pass, asking about things the student didn't know the first time through, and adding more detail to the structure of information built up on the first pass. The tutorial method as a whole reflected a strategy Norman (10) refers to as "web teaching", where the teacher first tries to establish a framework of basic knowledge and then fills in more and more detail on subsequent passes, much like a spider spinning a web.

The fourth aspect of the dialogues important to individualizing instruction is the way tutors correct student errors. When students make a confusion between two concepts, the better tutors try to provide distinguishing properties between the concepts for the

student. For example, when one student confused Ecuador and Colombia, the tutor pointed out that Ecuador is a much smaller country and that Colombia is connected to Panama. By providing distinguishing characteristics, the tutor is giving the individual the most relevant information for remembering the distinction in the future.

Tutorial Mode in SCHOLAR

These four aspects of the human tutoring strategy were developed in a mode called Tutor-SCHOLAR (9). Like the human tutor, Tutor-SCHOLAR selects topics in order of importance, and goes into depth on those topics the student knows something about. It starts out asking questions always probing deeper until the student doesn't know an answer or the time is used up for that topic. When the student misses a question, it presents a little related information for the student to assimilate before going onto the next topic. It allocates its time between a first pass and a review pass. On the review pass it skips over what the student knew earlier. However, it asks about everything the student missed, or that was presented earlier. Anything the student remembers from the earlier pass leads the discussion into more depth on the review pass.

When the student suggests an answer that differs from what is stored, Tutor-SCHOLAR tries to find a basis for rejecting the student's answer. This involves an inferential comparison between what the student says and what is stored. If it can find a

contradiction, the program points out the basis for rejecting the student's answer plus anything it knows about the relation of the student's answer to the object in question. In this way it tries to give the student a basis for remembering the distinction in the future.

Figure 2 shows a transcript between Tutor-SCHOLAR and a naive student who was testing out the system. It shows the kinds of mistakes SCHOLAR makes in language understanding. But it also shows the interactive flavor of SCHOLAR with a tutoring strategy directing the conversation.

Systematic Evaluation of Teaching Methods

Because the information in SCHOLAR is distinct from the procedures used for teaching that information, it is possible to vary the teaching strategy while holding the information constant. This makes it possible to compare different teaching methods in a systematic manner. To evaluate the effectiveness of Tutor-SCHOLAR we compared it to a teaching strategy derived from programmed instruction, which was implemented in a mode called Block-SCHOLAR. Block-SCHOLAR selects topics in the same order as Tutor-SCHOLAR, but unlike Tutor-SCHOLAR it first presents material and then asks questions about the material presented.

Three experiments were run to compare Tutor-SCHOLAR and Block-SCHOLAR (11). The comparison was made by teaching students the same domain of information with the two versions of SCHOLAR, and

LET'S BEGIN BY TALKING ABOUT THE LOCATION OF BRAZIL.   TRUE
OR FALSE? BRAZIL IS IN NORTHEASTERN SOUTH AMERICA.

THAT IS FALSE

NO. BRAZIL IS IN NORTHEASTERN SOUTH AMERICA.

THE EASTERN COAST OF BRAZIL IS THE   ATLANTIC.   THE   WESTERN
BORDERING COUNTRIES   ARE   PERU AND COLOMBIA.   THE SOUTHERN
BORDERING COUNTRIES ARE URUGUAY AND ARGENTINA.

> *Because the student missed the question, SCHOLAR presents
> some related information about the location of Brazil,
> and then goes on to the next topic.*

NEXT WE'RE GOING TO DISCUSS THE TOPOGRAPHY OF BRAZIL.   THE
PRINCIPAL RIVERS OF BRAZIL ARE----.

THE AMAZON AND THE ORINOCO.

THE AMAZON IS CORRECT. THE ORINOCO IS IN VENEZUELA AND
COLOMBIA, NOT IN BRAZIL.   THE   OTHER PRINCIPAL RIVERS OF
BRAZIL ARE THE TOCANTINS, THE SAO FRANCISCO, THE RIO NEGRO,
AND THE MADEIRA.

THE SAO FRANCISCO IS IN EASTERN BRAZIL.   IT IS IN THE
BRAZILIAN HIGHLANDS. THE RIO NEGRO IS IN THE AMAZON JUNGLE.
IT IS IN NORTHWESTERN BRAZIL.

> *Because the student knew about the Amazon, SCHOLAR takes
> it up as a topic for a while.*

CORRECT OR INCORRECT? THE AMAZON IS NAVIGABLE.

POSSIBLY YES

NO. THE AMAZON IS NAVIGABLE.

> *SCHOLAR misunderstood the "possibly" and treated the
> answer as wrong.*

THE AMAZON IS IN THE AMAZON JUNGLE. IT IS IN BRAZIL AND
PERU.   THE   PRINCIPAL TRIBUTARIES ARE THE RIO NEGRO, THE
MADEIRA, THE TOCANTINS, THE TAPAJOZ, THE PURUS, AND THE
JURUA.

WHAT IS THE LENGTH OF THE AMAZON?

Fig. 2 Dialogue with Tutor-SCHOLAR

IS THE AMAZON A LAKE?

I DON´T KNOW IF THE AMAZON IS A LAKE, BUT IT IS A RIVER.

> *SCHOLAR found that both lakes and rivers were water,*
> *so it couldn't reject the idea.*

NOW YOU CAN EITHER ANSWER MY QUESTION OR ASK A NEW ONE.

WHAT IS THE QUESTION?

WHAT IS THE LENGTH OF THE AMAZON?

5000 MILES

NO.   THE LENGTH OF THE AMAZON IS 4,000 MILES.

> *Having exhausted its allotted time to rivers, SCHOLAR*
> *went onto the next subtopic under topography of Brazil.*

THE PRINCIPAL MOUNTAINS OF BRAZIL ARE -----.

I AM NOT SURE, ARE THE ANDES IN BRAZIL?

IN WHAT YOU SAY, I DON´T RECOGNIZE THE WORDS "I", "AM" AND
"SURE".   TRY AGAIN.

ARE THE ANDES IN BRAZIL?

NO.   THE ANDES ARE NOT IN BRAZIL, BUT IN VENEZUELA,
COLOMBIA, ECUADOR, PERU, BOLIVIA, CHILE, AND ARGENTINA.

FEEL FREE TO ANSWER THE QUESTION YOU DIDN´T ANSWER BEFORE.

BRAZIL HAS NO MOUNTAINS.

NO. I DON´T KNOW WHAT YOU MEAN BY NO.   THE   CORRECT ANSWER
IS THE BRAZILIAN HIGHLANDS.

> *SCHOLAR didn't understand the answer, but proceeded quite*
> *reasonably.*

THE BRAZILIAN HIGHLANDS ARE IN SOUTHEASTERN BRAZIL.   THEY
ARE NEAR THE ATLANTIC.

WHAT IS THE PRINCIPAL PLATEAU OF BRAZIL?

> *The dialogue continued discussing different aspects of*
> *Brazil, but we will break here to the beginning of the*
> *second pass.*

Fig. 2 Dialogue with Tutor-SCHOLAR (cont)

NOW WE WILL GO OVER THE MATERIAL AGAIN. LET'S BEGIN BY TALKING ABOUT THE LOCATION OF BRAZIL.

PLEASE INDICATE IF THE FOLLOWING STATEMENT IS CORRECT OR INCORRECT: BRAZIL IS IN NORTHEASTERN SOUTH AMERICA.

TRUE

THAT'S RIGHT. BRAZIL IS IN NORTHEASTERN SOUTH AMERICA.

> *Because on this pass the student was correct, SCHOLAR asked questions about the information that was presented on the first pass.*

THE EASTERN COAST OF BRAZIL IS ----.

ATLANTIC SEA.

VERY GOOD. THE ATLANTIC IS CORRECT.

THE WESTERN BORDERING COUNTRIES OF BRAZIL ARE ----.

> *The dialogue continued through the second pass.*

Fig. 2 Dialogue with Tutor-SCHOLAR (cont)

measuring their learning by comparing pre-test and post-test scores on the material covered. In the three experiments, Block-SCHOLAR was systematically converged toward Tutor-SCHOLAR in order to pinpoint what aspects of teaching strategy affected student's learning. Tutor-SCHOLAR was significantly more effective in the first two experiments, and nonsignificantly in the third. The results indicated that the major factor affecting student's learning was the strategy that tutors use of reviewing the material in greater depth on a second pass. Allowing the students to ask questions, and the tutorial strategy of relating new material to the student's previous knowledge contributed a smaller amount to the greater effectiveness of Tutor-SCHOLAR.

The fact that SCHOLAR can be used to test particular aspects of teaching methods makes it potentially a valuable tool for educational research. The possibility of trying out single modifications in teaching strategy to see their effects on student's learning rate is unique. Human teachers of course can make such modifications in their own teaching strategies, but there is no way to control all the other factors that might vary as they change strategy. However, any specific version of SCHOLAR is a fixed system, and so an unbiased comparison can be made using any number of subjects. In this way the accumulation of systematic knowledge about teaching methods can occur.

TUTORING DIFFERENT TYPES OF KNOWLEDGE

One of the limitations of the original SCHOLAR was that it was restricted to teaching verbal facts about geography. The SCHOLAR system itself has been extended to encompass two other kinds of knowledge: visual knowledge about maps in the geography domain, and procedural knowledge about how to use a computer text-editing system called NLS (12). A related program called SOPHIE (13) tutors the diagnostic skills needed in electronic troubleshooting. In this section we will try to show the generality of this approach to CAI, and also some of the specific adaptions that have occurred in applying it to different domains of knowledge.

Tutoring Visual Knowledge

In order to explore the tutoring of visual information in an integrated manner with verbal information, we developed a Map-SCHOLAR System (14). The system can discuss with the student different maps that change dynamically according to the context of the discussion. To do this a graphic structure was created which parallels the structure in the semantic network. The elements in the map display therefore can be referred to either by their name, or by pointing to them, or both. Map-SCHOLAR both asks and answers map-related questions and provides relevant map information when the student makes a mistake. It has all the capabilities of Tutor-SCHOLAR with the addition of the map capabilities.

16

Figure 3 illustrates some of the variety of interactions possible with Map-SCHOLAR. The first panel illustrates how Map-SCHOLAR asks map-related questions, evaluates the student's answers, and corrects his mistakes. First, the dots indicating the location of the cities appear unlabelled on the map of Brazil and start blinking. Then SCHOLAR asks the student to name the blinking cities. When he responds, it tells him which answers were correct and prints their names on the map. Because he erroneously gave Lima as an answer, Map-SCHOLAR pointed out that Lima is in Peru and showed where Lima is on the map. Thus Map-Scholar gives both verbal and visual feedback to help the student learn the material.

The second panel shows how context affects the map display. The topic is the Amazon, which appears in this context on the map of Brazil, the overall topic. The question asks the student to point at the cities along the Amazon in the order named. The dots for the three cities first appear unlabelled, and only after the student has answered do the labels appear. Unlike the first map, this map shows only the cities on the Amazon and their relation to the Amazon. There is also expansion of detail so that a city like Belem, which was not important enough to be included among the cities of Brazil, is included with the subtopic of the Amazon. Because the map changes dynamically as the context changes and as the student interacts with the system, the students attention is focused on the relevant information exclusively and questions can be posed in a visual form not possible with a static map display.

NOW LET'S SEE WHAT YOU REMEMBER ABOUT THE CITIES OF BRAZIL.
PLEASE NAME THE PRINCIPAL CITIES OF BRAZIL NOW BLINKING
ON THE DISPLAY SCREEN.

MANAOS SAO PAULO AND LIMA

MANAOS AND SAO PAULO ARE CORRECT. LIMA IS IN PERU, NOT IN BRAZIL.
THE PRINCIPAL CITIES OF BRAZIL ARE BLINKING ON THE SCREEN.

RIO DE JANEIRO IS IN SOUTHEASTERN BRAZIL.
IT IS ON GUANABARA BAY AND THE ATLANTIC.

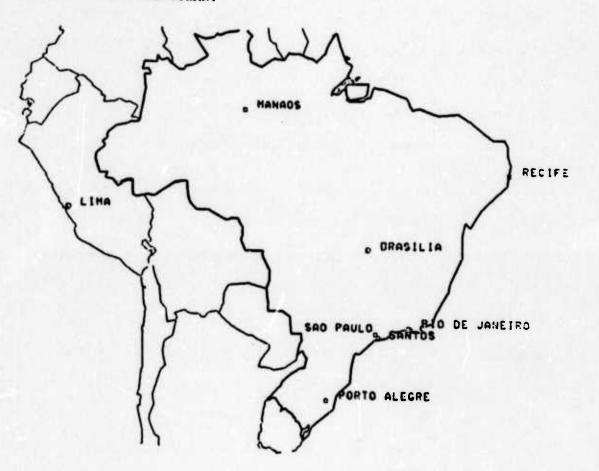BRASILIA IS IN CENTRAL BRAZIL.
IT IS FAR FROM THE OCEAN.

Fig. 3 Interactions with Map-SCHOLAR

18

PLEASE POINT TO THE FOLLOWING CITIES OF THE AMAZON
IN THE ORDER IN WHICH THEY APPEAR BELOW:
MANAOS, BELEM, AND IQUITOS
DON'T FORGET TO LET ME KNOW WHEN YOU ARE FINISHED,
PLEASE REMEMBER TO TYPE AN ASTERISK * FOLLOWED BY A
CARRIAGE RETURN.

**

MANAOS IS CORRECT. YOU DID NOT POINT AT BELEM AND IQUITOS CORRECTLY.
THE CITIES OF THE AMAZON ARE BLINKING ON THE SCREEN.

Fig. 3   Interactions with Map-SCHOLAR (cont)

LET'S BEGIN BY TALKING ABOUT THE TOPOGRAPHY OF VENEZUELA.
THE PRINCIPAL LAKE OF VENEZUELA IS BLINKING, PLEASE NAME IT.

I DON'T KNOW

THE PRINCIPAL LAKE OF VENEZUELA IS BLINKING ON THE SCREEN.

LAKE MARACAIBO IS IN NORTHWESTERN VENEZUELA.
THE PRINCIPAL MINERAL IS OIL.

IS THIS COUNTRY COLOMBIA?

YES, THAT'S COLOMBIA

WHERE IS THE CAPITAL OF VENEZUELA?

CARACAS IS BLINKING.

Fig. 3 Interactions with Map-SCHOLAR (cont)

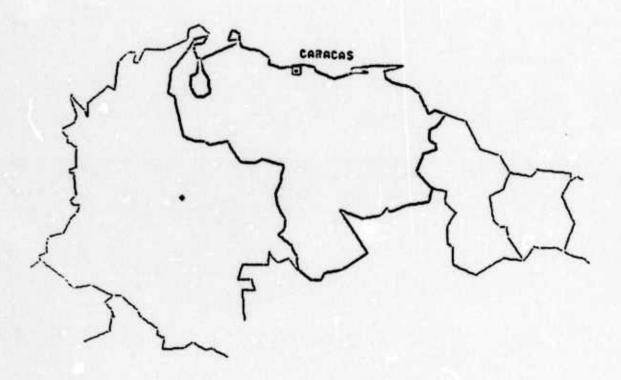The third panel shows the system's potential for tutorial interaction. When the student didn't know about Lake Maracaibo, Map-SCHOLAR showed it on the screen and added the related verbal information about the oil there. The two questions by the student illustrate some of the variety of questions the student can ask Map-SCHOLAR to clarify or amplify the information given (14). The first question involves both pointing at and naming Colombia to ask if that is the country near Lake Maracaibo. The second question shows another kind of map question. The subject verbally asks where the capital of Venezuela is, perhaps to find out how far away it is. SCHOLAR figures out semantically that the capital is Caracas, and then visually shows where Caracas is by blinking it (it is the double square). These examples illustrate some of the power for tutorial interaction that can be obtained by a close integration between semantic and visual knowledge.

There is now substantial evidence in the psychological literature (15) that people remember imageable material better than non-imageable material. Furthermore, if they use image strategies for remembering (16), they remember any given information better than if they use other strategies. Hence the fact that information is presented visually should make it more memorable.

As Collins & Quillian and Norman (10) argue, the best way to learn something is to relate it as much as possible to whatever information one already knows. Hence, even non-visual information like the fact that the Manaos has a tropical climate, will be

learned better, if one can see where Manaos is. This is true for two reasons. First, because when Manaos is located visually, it will be remembered better, and so facts that are related to it will also be remembered better. Second, if a student sees that Manaos is on the Amazon, then its climate can be related to any knowledge about the climate of the Amazon. Thus, information that seems not to depend on imagery should be remembered better in a visual context. For these reasons we expect the map facility, to have a substantial impact on how much students learn.

We have conducted an experimental evaluation of the map system using the comparative method described earlier (11). The test compared student's learning with Tutor-SCHOLAR using the map system vs. a labeled map vs. an unlabelled map. The experiment found an advantage of the map system over either of the static maps. We are using a technique called "backtrace analysis", which involves comparing the specific information each student learned with how that information was discussed, in order to pinpoint what aspects of the map system led to better learning by the students.

Tutoring Procedural Knowledge

NLS-SCHOLAR (17) is a prototype system to teach computer-naive people how to use the powerful NLS text-editing system (12). This teaching is accomplished by presenting a sequence of lessons. During each lesson the student may interact with the system by asking and answering questions, performing tasks which are posed by

the system, and performing tasks of his own choosing. Tasks are executed on an actual NLS system. Those tasks which have been posed are evaluated by the system, and the student is given encouragement, advice, and assistance.

NLS-SCHOLAR has been designed with the belief that procedural knowledge is best learned 'by doing'. It is an example (18) of a new kind of CAI system that integrates systematic teaching with actual practice, i.e., one which can keep a student under "intelligent" tutorial supervision while allowing him to try out what he learns on the very system he is learning about. Thus the system "knows" what the student is doing and can point out his mistakes, give specific help, show him how to do things and even do them for him.

NLS-SCHOLAR delivers a series of lessons designed for gradual understanding of NLS concepts and commands. Within these lessons, the system pauses to ask the student questions and to propose editing tasks for him to perform using NLS. A student's responses to questions and his performance of tasks are evaluated by the system and if he makes an error, the nature of his mistake is pointed out and appropriate action is taken. For example, if a question is answered unsatisfactorily, NLS-SCHOLAR proposes another question of the same kind. If a task is performed incorrectly, depending on the magnitude of the error, NLS-SCHOLAR either resets it for the student to try again, or asks him to proceed and try to fix his mistake, aided by the information NLS-SCHOLAR provides.

The user can formulate requests in relatively unconstrained English.   The requests can be questions about NLS concepts or about the state of his work, requests for help in doing a task, or even NLS commands expressed in English.   The system is "aware" of what the user is currently doing so that his requests for help can be answered within the context of the problem he is working on.   Thus NLS-SCHOLAR not only tells him "The general procedure is..." but also "In your case, what you should do is...".

NLS-SCHOLAR has the ability to use the NLS file a student is currently working on to show him how to perform editing actions. This gives the system much of the flavor of a human tutor, as if he were taking the student's place at the terminal and saying "Watch me do it for you".

The flavor of NLS-SCHOLAR is best conveyed by an annotated protocol, shown in Fig.   4, which was obtained on-line using a version of the system. In the course of a lesson, students learn how to change the contents of a menu by performing editing operations. The protocol starts at a point well along in the student's learning of NLS -- he has been told about NLS files, how to load them, print them, delete and insert statements, etc. He is about to be taught how to use the Substitute command to effect a change in the menu. Notice that what the student actually typed is underlined and that our annotations appear in italics.

The capabilities of NLS-SCHOLAR allow students to learn from explanation, learn by doing, and learn by asking questions. Their

24

```
BREAKFAST MENU

1   JUICE
    1A   ORANGE
    1B   GRAPEFRUIT
2   CEREAL
    2A   OATMEAL
         2A1   WITH RAISINS
    2B   CREAM OF WHEAT
    2C   CORN FLAKES
3   EGGS
    3A   SCRAMBLED
    3B   FRIED
         3B1   SUNNY-SIDE-UP
         3B2   OVER-EASY
    3C   BOILED
4   BEVERAGE
    4A   HOT CHOCOLATE
    4B   TEA
         4B1   WITH LEMON
         4B2   WITH SUGAR AND CREAM
    4C   COFFEE
```

SUBSTITUTE WORD IN STATEMENT

Before we end this lesson, I'd like to show you how to change the
content of statements which have already been inserted in your file.

You've learned how to correct errors by deleting an entire statement
and inserting a new one in its place. The Substitute command may be
used to change just a word or even a single character at a time. For
example, here's how you can change a word in statement 2A1 so that it
reads WITH DATES instead of WITH RAISINS.

```
EDIT C: (Su )bstitute C: (W )ord in  C: (S )tatement at  A: (2A1
)(<CR>);
 <New WORD> T: (DATES)(<CR>);
 <Old WORD> T: (RAISINS)(<CR>);
Finished? Y/N: (Y) OK: (<CR>);
Substitutions made: 1
EDIT C: (Do )ne OK: (<CR>);
```

*NLS-SCHOLAR performs a command for the student much as a
human tutor would by leaning over the student and typing
the command on the student's terminal. Parentheses enclose
what the "tutor" actually "types", separating it from the
echoes and responses printed by NLS.
<CR> is just a symbol to denote that the "tutor" depressed
the return key. The semicolon is NLS's echo to <CR>.*

Fig. 4 Dialogue with NLS-SCHOLAR

25

You can see that I've specified that I want to substitute one word in
statement 2A1 for another; the new word DATES is to be
substituted for the old word RAISINS.
(There is no restriction that the items substituted for
one another be of the same length.)

To do this I typed only 'Su W S 2A1 <CR>' to produce the first line; I
typed 'DATES <CR>' and 'RAISINS <CR>' to produce the second and third.
The remaining characters are all printed by NLS.

NLS asks 'Finished?' since more than one substitution can be made at a
time. If I had typed 'N' for no instead of 'Y' for yes, it would have
prompted for another <New WORD> and <Old WORD> pair.

After I typed 'Y' NLS asks for further confirmation with OK:
to which I responded with a <CR>.

When the command has been executed, NLS tells how many substitutions
have been made. If there had been two instances of RAISINS in
statement 2A1, both would have been changed to read DATES.

Statement 2A1 is now the current statement. Please print it.

Student typing is underlined.

TASK 41

EDIT C: /
  ==>WITH DAT
EDIT C: ↑
  2A OATMEAL

*The student confuses the \ command, which
prints the current statement, with the / command
that shows the current position of the control
marker. Realizing his mistake, he fumbles with
the other one-character commands he can
remember. First he tries ↑ which prints the
previous statement (2A), and makes it become
the current statement.*

EDIT C: __%

  2A1 WITH DATES
EDIT C: Done OK: __ ;

*Next the student depresses the LINE FEED
key which echoes as %, giving him the next
statement (2A1), making it current.*

You printed too many statements.   *The task is evaluated.*
This is what I wanted you to print:
  2A1 WITH DATES

        Please redo just the printing
part of this task. Your file is unchanged.

  *NLS-SCHOLAR is really determined to get the student to review
  how to print the current statement.*

Fig. 4 Dialoque with NLS-SCHOLAR (cont)

26

continuation of TASK 41

EDIT C: *What command prints the current statement?*

THE USER PRINTS THE CURRENT STATEMENT USING THE
\ COMMAND.

> *Failing to remember the appropriate one-character command,*
> *the student gives up and asks for help.*
> *Questions can be asked anytime it is the student's turn to type.*
> *NLS-SCHOLAR answers this question by retrieving information from*
> *its own internal, descriptive data base.*

continuation of TASK 41

EDIT C: \
 2A1 WITH DATES
EDIT C: Done OK:_;

QUESTIONS?

**_

> *The student chooses to ask no questions at this time.*

Try the Substitute command yourself by substituting the word
OAT for CORN in the statement containing "CORN FLAKES".

TASK 42

EDIT C: *What statement contains "CORN"?*

THE STATEMENT CONTAINING "CORN" IS:
 2C CORN FLAKES

> *The student wants to know what the correct statement number is*
> *for him to type. The system obtains the answer by actually*
> *looking into the student's work space -- the answer is not stored*
> *(it cannot be!) in the internal data base.*
> *The system interprets correctly most of the common*
> *paraphrases of this type of question, e.g.*
> *"What is the statement containing ..."*
> *"What statement contains ..."*
> *"Where is ...", etc.*

Fig. 4   Dialogue with NLS-SCHOLAR (cont)

continuation of TASK 42

```
EDIT C: Substitute C: Word in  C: Statement at  A: 2C;
 <New WORD> T: CORN;
 <Old WORD> T: OAT;
Finished? Y/N: Y  OK:_;
Substitutions made: 0
EDIT C: Done OK:_;
```

*The student switches around the new and the old words,
so no change is made.*

You should have altered the file
so that this part of it
    2B CREAM OF WHEAT
    2C CORN FLAKES
  3 . . .

*NLS-SCHOLAR shows the student
what the result of this command
should have been if he had
performed the task correctly
as opposed to the result which
he obtained. With this
information, the student
can "debug" his performance
of the task.*

would look like this:
    2B CREAM OF WHEAT
    2C OAT FLAKES
  3 . . .

But you did not alter the file.

      Please try to repair your
mistake. The file is as you left it.

continuation of TASK 42

EDIT C: *How does one do this?*

*The student is lost and
asks for help.*

THE USER TYPES 'SUBSTITUTE', FOLLOWED BY 'WORD',
'STATEMENT', '2C', <CR>, 'OAT', <CR>,
'CORN', <CR>, 'Y', <CR>, 'DONE', AND <CR>.

*NLS-SCHOLAR can show him what
to do, as well as...*

continuation of TASK 42

EDIT C: *Please do it for me*

...actually do it for him!

```
EDIT C: (Substitute ) C: (Word ) in C: (Statement ) at A: (2C)(<CR>);
 <New WORD> T: (OAT)(<CR>);
 <Old WORD> T: (CORN)(<CR>);
Finished? Y/N: (Y) OK: (<CR>);
Substitutions made: 1
EDIT C: (Done ) OK: (<CR>);
```

Fig. 4 Dialogue with NLS-SCHOLAR (cont)

tight integration within a working environment makes NLS-SCHOLAR a powerful tutorial CAI system.

Tutoring Diagnostic Skills

SOPHIE (13) reflects an attempt to extend Carbonell's notion of mixed-initiative CAI for the purpose of encouraging a wider range of student initiatives. Unlike previous tutorial systems which attempt to mimic the roles of a human teacher, SOPHIE tries to create a "reactive" environment in which the student learns by trying out his ideas rather than by instruction. To this end, SOPHIE incorporates a "strong" model of its knowledge domain along with numerous heuristic strategies for answering a student's questions, providing him with critiques of his current solution paths, and generating alternative theories to his current hypotheses. In essence, SOPHIE enables a student to have a one-to-one relationship with an "expert" who helps the student create, experiment with, and debug his own ideas.

SOPHIE's expertise is derived from an efficient and powerful inferencing scheme that uses multiple representations of knowledge including (a) simulation models of the domain (b) procedural specialists which contain logical skills and heuristic strategies for using these models, and (c) semantic nets for encoding time-invariant factual knowledge. The power and generality of SOPHIE stems, in part, from the synergism obtained by focusing the diverse capabilities of the procedural specialists on the

"intelligent" manipulation, execution, and interpretation of its simulation models.

In the basic scenario, SOPHIE acts as an electronics lab instructor who helps the student transform his classroom knowledge of electronics into an experiential, intuitive knowledge of its meaning and application. It does this ьy interacting with the student while he is debugging a malfunctioning piece of equipment (19). The student can perform any sequence of measurements, ask either specific questions about the implications of these measurements for more general hypothetical questions, and even ask for advice about what to consider next, given what he has discovered thus far. At any time SOPHIE may encourage the student to make a guess as to what he thinks might be wrong given the measurements he has made thus far. If he does, SOPHIE will evaluate his hypothesis by taking into consideration all the information he should have been able to derive from his current set of measurements. If any of this information is logically contradicted by the hypothesis, SOPHIE identifies and explains these contradictions. Likewise SOPHIE can judge the merits of any particular measurement with respect to the prior sequence of measurements he has made. For example, his new measurement may be logically redundant in the sense that no new information can possibly be derived from it (an extremely complex task to determine). SOPHIE can also decide if this measurement performs a reasonable split of the hypothesis space of possible faults which have not yet been ruled out by prior measurements.

It should be noted that the scenario contains quite a variety of logical tasks (i.e., hypothesis evaluation, hypothesis formation, redundancy checking, hypothetical question answering) each one of which requires a substantial amount of deep logical inferencing. One of the basic challenges in constructing SOPHIE was creating an inference system which could perform this wide range of tasks efficiently (so that it could be used in real time) and at the same time have it be robust in the sense of handling all realistic queries.

Because SOPHIE was designed as an environment in which students could create and articulate ideas, it was necessary to have a powerful natural language processor to communicate with students. A student will become frustrated if he has to try several ways of expressing an idea to get a response. In addition he will become bored if there is a long delay (say 10 secs) before the system replies. And because students begin to assume the system shares their "world-view", SOPHIE must cope with contextually-dependent references, deletions, and ellipses. SOPHIE's natural language processor is based upon a "semantic grammar" technique, in which concepts like "measurement" or "circuit element" trigger expectations about what things should appear in the student's input. SOPHIE has demonstrated that natural language processing has advanced far enough to deal with these three kinds of difficulties well enough to build friendly, but sophisticated tutorial systems.

THE FUTURE OF INTELLIGENT CAI

The thrust of this paper has been to show what kind of capabilities are now available for building genuinely intelligent CAI systems. The domain of such systems is virtually unlimited; it is not restricted, for example, to drill and practice or mathematics. The language capabilities of current systems are not equal to those of a human, nor will they be in the foreseeable future, but they are good enough to sustain practical systems.

The Plato system (6) has shown that it is possible to build both interesting and cost-effective CAI systems in a time-shared computer environment. They have accomplished this by using a variety of teaching techniques: the Socratic method, generative CAI, games and simulations, programmed instruction, etc. Intelligent CAI is an attempt to go beyond the technology in the Plato system to explore how to build greater intelligence into tutorial systems, while at the same time utilizing many of the educational techniques employed so successfully in Plato.

Intelligent CAI systems are now both costly to build (above $100,000) and to use (about $10-$20 per hour). But, the cost of computing continues to decrease while teacher's salaries are rising. Hence the cost of running such systems should be competitive in comparison to the cost of human tutoring within a short time, especially where there are few skilled teachers available, as with teaching computer text-editing. The effective cost of building such systems depends on how much they are used. If they are used

heavily, then the large cost of building them will be worth the investment; otherwise not. It is at least possible that one of the current systems will be used enough to justify the development expense, though they were built only as prototype systems. The test though will be the development of such a system for a school setting where large numbers of people are being taught.

The payoff in intelligent CAI comes from personalizing the learning process. Personalization is effective in many ways: by forcing the student to participate in learning; by teaching at the level of his individual knowledge; by providing a setting where the student can try out his own ideas and make mistakes; by freeing the student from peer pressure; by addressing the student's individual confusions, etc. These advantages make it worthwhile to give intelligent CAI a serious trial.

References and Notes

1. J.R. Carbonell, *IEEE Trans. Man-Mach. Sys.* MMS-11, 190 (1970).

2. G.L. Bryan, *Comp. & Automat.* 18, No. 3, 1 (1969).

3. S. Papert, *Prog. Learn. & Ed. Tech.* 9, No. 5, (1972). I.P. Goldstein, *AITR-294*, (AI Lab., MIT, Cambridge, Mass., 1974). W. Feurzeig & G. Lukas, *Ed. Tech.* 12, No. 3, 39 (1972).

4. C. Resnick, *Computational of Models of Learners for Computer Assisted Learning*, Unpublished Doctoral Dissertation. (U. of Illinois, Urbana, Ill., 1974).

5. J. Swets & W. Feurzeig, *Science* 150, 572 (1965).

6. D.L. Bitzer & R.L. Johnson, *Proc. of IEEE* 59, 960 (1971) D.L. Bitzer, B.A. Sherwood, P. Tenczar, *CERL Rpt. X-37*, (Computer-based Education Research Laboratory, U. of Illinois, Urbana, Ill., 1972).

7. L. Uhr, *Proc. 24th ACM National Conf.* (New York, 1969) p. 125. W.R. Uttal, T. Pasick, M. Rogers, and R. Hieronymus. *Comm.* 243. (Mental Health Research Inst., Ann Arbor, Mich.,

1969).   J.D.   Wexler,   IEEE   Trans.   Man-Mach.   Syst., MMS-11,
181, 1970.

8. M.R. Quillian in Semantic Information Processing, M.   Minsky
    Ed.   (MIT Press, Cambridge, Mass., 1968) p. 216. R. Simmons
    in Computer Models of Thought and Language, R.C.   Schank & K.
    Colby, Eds. (Freeman, San Francisco, 1973) p. 63.

9. A.   Collins, E.H.   Warnock, and J.J.   Passafiume in The
    Psychology of Learning and Motivation, Vol. 9, G.H. Bower, Ed.
    (Academic Press, New York, 1975).

10. A. Collins and M.R. Quillian in Cognition in Learning and
    Memory, L.W. Gregg, Ed., (Wiley, New York, 1972) p.   117. D.A.
    Norman in Contemporary Issues in Cognitive Psychology: The
    Loyola Symposium, R.L.   Solso, Ed. (Halsted Press, New York,
    1973).

11. A. Collins, BBN Report No. 2885. (Bolt Beranek & Newman,
    Cambridge, Mass., 1974).

12. NLS, the On Line System, is a sophisticated modular system which
    is being used increasingly as an aid in writing, re-organizing,
    indexing, publishing, and disseminating information of all
    kinds. It was developed by Douglas Engelbart and his co-workers
    at the Augmentation Research Center of the Stanford Research

Institute.

13. J.S.   Brown   and   R.R.   Burton   in   <u>Representation</u>   <u>and</u>
    <u>Understanding</u>,   D.G.   Bobrow   and   A.   Collins,   Eds.   (Academic
    Press, New York, 1975) p.   311.   J.S.   Brown, R.R.   Burton,   and
    A.G.   Bell, <u>Intl.</u>   <u>J.</u>   <u>Man-Mach.</u>   <u>Stud.</u>, in press (1975).

14. A.   Collins and E.H.   Warnock, <u>BBN</u>   <u>Report</u>   <u>No.</u>   <u>2833</u> (Bolt
    Beranek & Newman, Cambridge, Mass., 1974).

15. A.   Paivio, <u>Imagery</u> <u>and</u> <u>Verbal</u> <u>Processes</u>.   (Holt, Rinehart,   and
    Winston, New York, 1971).

16. D.A.   Norman, <u>Memory</u> <u>and</u> <u>Attention</u>.   (Wiley,   New   York,   1969).
    G.H.   Bower   in   <u>Cognition</u> <u>in</u> <u>Learning</u> <u>and</u> <u>Memory</u>, L.W.   Gregg,
    Ed.   (Wiley, New York, 1972) p.   51.

17. M.C.   Grignetti, C.   Hausmann,   and   L.   Gould   in   <u>Proc.</u>   <u>of</u>
    <u>National</u> <u>Computer</u> <u>Conf.</u>   (Anaheim, Calif., 1975) p.   775.   M.C.
    Grignetti, L.   Gould, A.G.   Bell, C.L.   Hausmann,   G.   Harris,
    and   J.J.   Passafiume   <u>BBN</u>   <u>Report</u>   <u>No.</u>   <u>2969</u>.   (Bolt Beranek &
    Newman Inc., Cambridge, Mass., 1974).

18. Other examples are reported in A.   Barr,   M.   Beard,   and   R.C.
    Atkinson,   <u>TR-228</u>   (Psychology and Education Series, Stanford U.
    Stanford, Calif., 1974) and in D.R.   Gentner, M.R.   Wallen   and

P.L.   Miller, A Computer-based System for Studies in Learning, (Center for Human Information Processing, UCSD, LaJolla, Calif., 1974).

19. Although the domain of knowledge under consideration is electronics, the reasoning and linguistic paradigms underlying SOPHIE are applicable to many domains outside of electronics.